UNIT 7

THE PROBLEM OF INDUCTION

1. Introduction

We have now seen how the rationalists and the empiricists answer the skeptic. We turn now to the third of the British Empiricists, David Hume (1711-1776), who discovered that the empiricist view itself leads to a serious skepticism. Very roughly, Hume questioned how the empiricists can claim to have any knowledge of the unobserved – e.g. of the future. Nowadays, we seem to think that we do know (at least sometimes) what will happen in the future because we have scientific knowledge. We know, for example, that the Payne Center will not collapse in and of itself tomorrow or next month because it has been well designed and the laws of physics, together with our knowledge of sufficiently many of them, underwrite our confidence that it is quite safe to enter the building. Hume showed that a pure empiricist has no reasons for such confidence.

2. Observational and Inductive Knowledge

Empiricists claim that knowledge derives from experience. But contrast the following three claims:

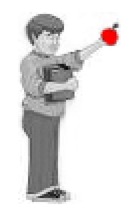

(1) This apple is red.
(2) This apple will fall when I drop it.
(3) All objects dropped near the surface of the Earth will fall down.

Claim (1) we can be quite confident about because it is justified by observation (of course, after all the traditional skeptical worries have been quieted but let's assume that they have). Claim (1) is an observational claim. Its truth (or falsehood) can be determined by observation alone.

Claim (2), on the other hand, is not an observational claim. In making it (now) we are predicting what will happen in the near future. Claim (2) might appear so obvious to you that you might doubt that it should not be put on a par with claim (1). But consider other claims like (2):

(2a) This apple will taste sour when I eat it.
(2b) This apple will taste sweet when I eat it.
(2c) This apple will taste like chocolate when I eat it.

Are we justified in affirming or rejecting any of these three claims just in virtue of our observing this apple? It seems clear that we are not. We would most likely reject (2c). Why? Well, because we know that *no apple tastes like chocolate!* But note that in thus rejecting claim (2c) we are appealing to *general* knowledge that no apple tastes like chocolate. To the extent that we might want to affirm (2a), for example, the only reason for us to do so would be our knowledge that this apple is of a kind that always tastes sour. In doing so, we are again appealing to *general* knowledge. We are not relying on observation alone. The same is true of claim (2). To the extent that we are confident that this apple will fall when the boy drops it, we are appealing to *general* knowledge of the sort capture in claim (3), for example. Thus claim (2) is not an observational claim. Its truth (or falsehood) cannot be determined by observation alone.

Having said that much, it should be clear that claim (3), viz. that all objects dropped near the surface of the Earth fall down toward the Earth (and not away from it, for example), is not an observational claim either. For one, we have not observed *all* objects!

Empiricists are, of course, aware of the fact that we do not (now) have observational knowledge of the future, for example. But, they claim, this does not mean that we cannot have any knowledge of the future. We can be justified in making claims about the future by induction. According to the empiricists, general (empirical) knowledge is always inductive knowledge.

3. Induction

3.1. What Is Induction?

Suppose that we want to know what color ravens are. How would we find out? Well, surely not by using philosophical proofs and trying to prove (or disprove) the thesis that all ravens are pink, for example. Clearly, we should just go out and see. And what do we see? We see one raven – it's black. We see another raven – it's black. We see yet another – it's black too. The next hundred ravens are all black. The next thousand ravens are all black. At this point (and perhaps a little earlier), we might be quite exhausted but we should slowly get the hunch that we have quite good reasons to believe that… all ravens are black.

Of course, we cannot still rule out the possibility that raven_{1110} will not be black. But we really do have *good* reason for thinking that raven_{1110} will be black. Moreover, our good reason for thinking that the next raven we see will be black is – we want to say – an *empirical* reason. It is not as if we have formed the hypothesis that ravens are black sitting in an armchair. We have formed it on the basis of experience.

The form of inference that allows us to infer something about the whole population of ravens (or Xs in general) on the basis of observing a sample of ravens (or Xs) is called induction. Here is its general form:

Raven_1 is black	X_1 has property P
Raven_2 is black	X_2 has property P
Raven_3 is black	X_3 has property P
⋮	⋮
So, all ravens are black.	So, all *Xs* have property P

Inductive Inferences Are Fallible Inferences

We will be talking about the problem of induction quite soon. However, it is important for you not to misunderstand the problem and not to think that you already see some problem with induction.

At this point, you certainly see that induction is a fallible inference – the truth of its premises does *not* guarantee the truth of the conclusion. (An infallible, or demonstrative, inference is an inference where the truth of the conclusion is guaranteed by the truth of the premises. A fallible, or non-demonstrative, inference is an inference where the truth of the conclusion is not guaranteed by the truth of the premises). There have been numerous examples to demonstrate it. To keep in line with the raven type of example, we can recall the conclusion that the

Europeans have reached about swans. All the (adult) swans they have observed were white, so they concluded (using induction) that all swans are white. Further exploration showed that the Australian swans are black and the South-American swans are black and white (black-necked swans are white except for head and neck while the Coscoroba swans have white feathers with black tips). So, in the end, the original inductive inference turned out to be faulty.

But the fact that inductive conclusions might turn out to be false is a feature of, not a problem with, induction. There are no guarantees in the empirical knowledge unlike in formal (deductive) knowledge.

Inductive Inferences Are Reasonable Inferences

You might wonder why we should want induction if it is fallible. Well, this is a good question. But in fact, go back to the ravens example. If you really think about, does it not seem reasonable to draw the conclusion? You see raven$_1$, which is black. You see raven$_2$, which is black. You see raven$_3$, which is black. And all the next hundred ravens are black too. What are you going to conclude? That all ravens are pink? No. It seems only *reasonable* to conclude that all ravens are black *on the basis* of the evidence thus far.

And this surely seems what happened in the case of conclusion reached by the European birdwatchers. All the swans they could observe were white, which made it *reasonable* for them to conclude that all swans are white *on the basis* of the evidence they had. Of course, they had to reject the conclusion once new evidence came in. With the new evidence in place, it was reasonable for them (and us) to conclude (by induction, again) that (a) all swans are either white or black or white and black and (b) all Europeans swans are white.

Moreover, the more inductive inference you gather the more reasonable is it to hold the inductive conclusion. Inductive inferences (unlike deductive inferences) have a probabilistic character. If you have seen 10 ravens and they have all been black you have some reason to believe that all ravens are black. If you have seen 100 ravens and they have all been black, it is now more probable that the inductive conclusion that all ravens are black is true. And if you have seen 10000 ravens and they have all been black, the conclusion gains even more probability. But note that in an inductive inference you will *never* gain *certainty*.

So, while the truth of the inductive premises *does not guarantee* the truth of the inductive conclusion, the truth of the inductive premises *raises the probability* of the truth of the inductive conclusion and thus makes it more *reasonable* for us to accept the inductive conclusion.

General Form of Inductive Inferences

We have already talked about the general form of inductive inferences, but it will pay to present it in an abbreviated form:

All observed *X*s have property *P*

All *X*s have property *P*

(Note that the conclusion like is double to mark the inductive nature of the inference; in deductive inferences a single line is used to separate the premises from the conclusion.) It is easy to see that all our examples fit this form:

All observed ravens have been black.

All ravens are black.

All observed swans have been white.

All swans are white.

All observed Europeans swans have been white.

All European swans are white.

All observed swans have been white, black or white with black markings.

All swans are white, black or white with black markings.

But the inferences are, of course, not limited to the examples we have seen. To the extent that science offers us general knowledge, it is based on induction.

All observed physical objects obey the law of gravity.

All physical objects obey the law of gravity.

Our ordinary knowledge is also based on induction:

The sun has risen in the past.

The sun will continue to rise in the future.

3.2. Two Important Properties of Inductive Inferences

We have already said that inductive inferences are fallible (not certain – they never guarantee the truth of the conclusion), yet reasonable because they raise the probability of the truth of the conclusion. There are two further properties of inductive inferences that we need to discuss and that will slowly move us toward posing the problem of induction. First, inductive inferences are ampliative. Second, inductive inferences rely on a missing premise.

Inductive Inferences Are Ampliative

An inference is ampliative when its conclusion carries more information than do the premises. All inductive inferences are ampliative. To see it, consider the inference about the ravens:

All observed ravens have been black.

All ravens are black.

Why does the conclusion contain more information than the premise? Because the premise carries information only about the observed ravens, while the conclusion carries information about the observed ravens as well as about unobserved ravens. This will be the case for every inductive inference. The general form of inductive inference:

All observed *X*s have property *P*

All *X*s have property *P*

can be also represented thus

All observed *X*s have property *P*

All (observed and unobserved) *X*s have property *P*

The latter representation makes it very clear why the conclusion contains more information than the premise.

For those of you who know a little about deductive logic, I should add that all deductive inferences are not ampliative – the conclusion in a deductive inference never carries more information than the premises. The deductive conclusion can carry the same amount of information as the premises or it can contain less information than the premises. This is in fact what makes such inferences infallible. They never go beyond what was contained in the premises.

All Inductive Inferences Rely on a Missing Premise

The ampliative nature of induction underscores the fact that in an inductive inference we are making a leap – from what we have direct observational knowledge about to what we do not have direct observational knowledge about but – what we reasonably suppose will be like what we do have direct observational knowledge about.

It was David Hume who was the first to note that in making this leap we rely on a tacit premise that "the future will resemble the past" or that nature is uniform. Let me first say what the content of the tacit premise is and then why it is a premise of all inductive inferences (though it is not usually explicitly stated).

What does Hume mean when he says that the future will resemble the past? He cannot, of course, mean that we will all stay young for ever, that the world in the future will be like the world of the past. We all know that changes occur. What Hume means is that the regularities that govern the world stay the same. People may change but the pattern of the changes is the same – we all get older. An object dropped toward the Earth may *change* its position but the law governing the fall of the object does not change. And so on and so forth. It is sometimes said that the premise is better captured by the claim that nature is uniform, or that natural laws do not change.

Why does Hume believe that all inductive inferences rely on the assumption that nature is uniform? Hume thinks that only a person who believes that nature is uniform has any reason for making the inductive inference. Consider a person, Albert, who believes that nature is not uniform, that natural laws change. No matter how many observations to the effect that all observed physical objects have obeyed the law of gravity we present to Albert, he will not have any reason for thinking that all physical objects obey the law of gravity – this is because he believes that natural laws change. If we say to Albert "Look! All observed ravens have been black", he will have no reasoning for drawing the conclusion we draw, viz. "All ravens are black". This shows that the premise that natural laws do not change, that nature is uniform, is necessary for us to make an inductive inference. Someone who does not believe that nature is uniform has no reason to draw the inductive conclusion. Only a person who does believe that nature is uniform has a reason to draw the inductive conclusion.

In other words, to the extent that we can claim to know that all unobserved future ravens will be black, we must know it indirectly, relying on the fact that we know that all observed past ravens have been black *and* our belief that the future will be like the past. It is because we think that nature is uniform (that what happens in the future is in many respects like what happens in the past) that it so much as makes sense for us to take what happened in the past as a *ground* for believing that the same will hold in the future. It is because we think that nature is uniform that we take the fact that the law of gravitation has worked so far to constitute a good reason to believe that it will continue working in the future.

In other words, all inductive inferences rely on a premise that is often not explicitly stated (hence, the missing premise):

All observed *X*s have property *P*

Future resembles the past (i.e. nature is uniform)

All (observed and unobserved) *X*s have property *P*

4. The Problem of Induction

Now that you have learned a little about inductive inferences, you are ready for the encounter with the problem of induction. The problem of induction is a problem for the empiricists who believe that all knowledge must be justified empirically. We have also seen that induction is very important for the empiricists because it is by virtue of applying induction that they obtain general knowledge, i.e. knowledge that goes beyond the strictly observational. Induction thus promises to give general but still empirical knowledge. Hume discovered that (and in a moment we will see how) induction cannot be justified empirically.

4.1. Hume's Discovery: Induction Cannot Be Justified Empirically

We are dealing with an inductive inference, which we know has the form:

(O) All observed *X*s have property *P*

(M) Future resembles the past (i.e. nature is uniform)

(I) All (observed and unobserved) *X*s have property *P*

To the extent that the empiricists want to accept conclusions of inductive inferences as knowledge, they claim that inductive conclusions of the form (I) are *justified empirically*. This certainly seems like a reasonable assumption, as we have seen. What sort of justification do we have for holding that all ravens are black? Surely, (we think that) we are empirically justified in holding that all ravens are black – we draw this conclusion from empirical evidence.

However, Hume reminds us that a conclusion of an inference will be indeed justified empirically as long as the premises of the inference are justified empirically. In our case, the inductive conclusion (I) will be justified empirically as long as (1) the premise (O) is justified empirically and (2) the premise (M) is justified empirically. Let's check and see.

Is premise (O) justified empirically? The answer is surely positive. The claims that all observed ravens have been black, that all observed swans have been black are based purely on observation. They are thus empirically justified. In fact, premise (O) is sometimes just called the observational premise.

Is premise (M) justified empirically? Do we have empirical reasons for thinking that the future resembles the past, or that nature is uniform? The most natural answer is:

(n) Yes, we know that nature is uniform from experience – after all nature was uniform in the past.

Hume discovered that there is something really fishy about this response. If you think about it the very same problem of justification arises. We may know that nature was uniform in the past. But *on what grounds* do we think that this will continue to be the case in the future? This is the very problem that we started out answering. [You can already sense that some circularity is involved in this attempt to justify induction inductively, but here is a way of bringing the point to light more formally:]

Here is a way of showing more sharply that we are involved in circular reasoning when we try to offer an inductive justification of (M), i.e. when we try to justify (M) by appealing to experience. The structure of the reasoning in (n) is this:

(O^M) Nature was uniform in the past.

(M) Nature is uniform

So far so good, this is an instance of inductive reasoning. We infer something very general from what we know about the past. But recall that Hume has discovered that all inductive inferences carry with it an implicit premise: viz. that nature is uniform. (If we did not believe that the future will be like the past, we would have no grounds for drawing the inductive conclusion.) So, we need to supplement the inductive inference from (O^M) to (M) with the missing premise:

(O^M) Nature was uniform in the past.

(M) Nature is uniform.

(M) Nature is uniform

Now, however, it is apparent that the conclusion follows only because it is already assumed. The reasoning is circular.

The problem of induction consists in the fact that we cannot offer an empirical justification of the conclusions of inductive inferences in general. An attempt to show that inductive conclusions can be justified empirically ends in a vicious circle.

To the extent then that we think that scientific theories and laws are results of induction or are inductively supported by evidence, the problem of induction shows that ultimately we cannot empirically justify scientific theories and laws. Note also that the trouble-maker is the claim that nature is uniform. It is because we are at a loss trying to justify it that the problem arises. If we could justify (M), all else would fall in place.

4.2. The Problem of Induction as a Problem for Empiricism

It is important to emphasize that the problem of induction is a problem for empiricism. It is only someone who puts forth the requirement that all knowledge must be justified empirically that falls prey to the problem of induction.

Yet, it is also rather hard to see how one would want not to be an empiricist about the laws of science, say. If you ask a scientist on what grounds they hold their empirical laws, few would want to say anything else than what you would expect from an empiricist. In fact, it is the appeal to experience alone – not assumption, not reason, not faith – that is considered to be so unique in science.

I was once presenting the problem of induction to an audience that included some mathematicians. One of the mathematics professors toward the end raised his arm and asked the question "Can't we just assume that nature is uniform?" Of course, we can. But the problem is then that the whole body of our knowledge depends on this assumption. All empiricists (and most of us are empiricists about scientific knowledge at any rate) find it extraordinary that all general knowledge be based on an empirically unjustified assumption.

☯ Can we avoid the problem of induction if, instead of inferring that all Xs have property P, we draw the weaker inference that most Xs have property P?

4.3. Summary

The problem of induction can be summarized by means of three mutually incompatible propositions, all of which have been endorsed by the empiricists.

All knowledge must be justified empirically.

Some knowledge (in particular, all general knowledge) relies on induction.

Induction cannot be justified empirically.

5. Solving the Problem of Induction

There have been numerous attempts to solve the problem of induction, though most of them rely on the agreement with Hume that induction cannot be justified empirically. One kind of response was to argue that although one cannot justify induction empirically, it can be justified in another fashion (we will discuss one such argument, viz. Reichenbach's pragmatic justification of induction). This kind of response involves a partial abandonment of empiricism. Another kind of response has been to deny that general knowledge relies on induction. This response was proposed by Karl Popper, but to fully understand it you would need to know more about philosophy of science. I strongly recommend that you look at the *Logic of Scientific Discovery* (especially the first two chapters) if you are at all interested in this topic. Hume has proposed his own "solution" to the problem, but it is hard to see in what sense it is a solution to the problem.

5.1. A Pragmatic Justification of the Problem of Induction

Hans Reichenbach (1891-1953), a German philosopher, one of the prominent representatives of logical positivism, proposed that while we cannot offer empirical justification of induction, we can justify it pragmatically.

We are asking the question whether it is rational to use induction or not to use it. The relevant question about the world that would make a difference is: "Is Nature uniform?" We do not know the answer to this question, so we need to consider two possibilities: (a) nature is uniform, (b) nature is not uniform. In other words, we are consider the following four possibilities:

Choices: \ The world:	Nature is uniform	Nature is not uniform	expected payoffs:
Use induction	(a)	(b)	
Don't use induction	(c)	(d)	

What we are after is deciding whether we can be reasonably successful in testing scientific laws. It is easiest to see the point by focusing on the process of inducing generalizations (All ravens are black) from observation (This raven is black, That raven is black, etc.).

(b) If nature is not uniform and we use induction, we will be failing much of the time. Imagine that non-uniformity of nature is revealed in the fact that ravens change color randomly. So when you've assembled all the ravens in the sample, and all of them have been black, by induction you will infer that all future ravens will likewise be black. But if nature is not uniform, your conclusion will be a bad one, for ravens will change color.

(d) If nature is not uniform and we do not use induction, we will fail because we are not even trying to succeed.

(c) Likewise, even if nature is uniform and we don't use induction, we will fail (for we are not even trying.

(a) But, if nature is uniform and we use induction, while we will not succeed all the time (inductive inferences are infallible), we will succeed much of the time and we will succeed in the long run (with ever new evidence coming in).

Let's fill in the table:

Choices: \ The world:	Nature is uniform	Nature is not uniform	expected payoffs:
Use induction	(a) long-run success	(b) failure	Chance of success
Don't use induction	(c) failure	(d) failure	Guaranteed failure

If we now consider our two choices: whether to use induction or not to use it, we see that we are better off using it, because, if we use induction we at least have a chance of success, while if we do not use it, we are guaranteed a failure.

Note that this pragmatic solution to the problem of induction does not undermine Hume's insight that induction cannot be justified empirically. Reichenbach agrees with Hume in this respect. But he also thinks that it is still rational for us to use induction.

> ☯ Recall Pascal's pragmatic argument (discussed in Unit 2). Do you think that the fact that a pragmatic argument can be used both in justifying the foundations of scientific and religious knowledge is of significance? If so, what does it show?

5.2. Hume's "Solution" to the Problem of Induction

Hume believes that the foundation of inductive reasoning (reasoning from experience) is custom. We simply do make such inferences as a matter of habit. In fact, the habit is so well entrenched that we cannot possibly think of proceeding without it.

Although Hume advances this observation as a "solution" it is hard to see it as such. The *fact* that we *do* make inductive inferences does not establish that we *should* make inductive inferences, in other words, it does not justify inductive inferences. The fact that the use of induction is well entrenched does nothing to *justify* the belief that nature is uniform. The fact can only be used to show that we *do* hold the belief that nature is uniform.

Hume's solution seems to be just a version of the is-ought fallacy. Ironically, it was Hume who identified the is-ought fallacy in ethics. Hume objected that some moral theorists argued what *ought* to be the case on the basis of what *is* the case. Here, he seems to be doing the same thing.

6. Conclusion

The problem of induction is one of the central problems in epistemology but especially in the philosophy of science. All philosophers of science have addressed Hume's worry in some way or another. It challenges all empiricists to explain exactly what reason they have for believing anything about the future.